

09/29/00



10-Q-00

A

ASSISTANT COMMISSIONER FOR PATENTS
PATENT APPLICATION
Washington, D.C. 20231

CASE DOCKET NO. YOR920000465US1
Date: September 29, 2000

10929 U.S. PTO
09/29/00
09/676425



Transmitted herewith for filing is the Patent Application of:

Inventors: Tova ROTH et al.

For: NET ZEROING FOR EFFICIENT PARTITION AND DISTRIBUTION
Enclosed are:

☒ 7 Sheets of Informal Drawings.

☐ An assignment of the invention to International Business Machines Corporation, Armonk, New York 10504.

☐ A certified copy of a _____ application.

☒ Declaration and Power of Attorney is attached to the application.

☐ Associate Power of Attorney.

☐ Information Disclosure Statement with form PTO-1449 with references attached.

The filing fee has been calculated as shown below:

	(Col. 1)	(Col. 2)
FOR:	NO. FILED	NO. EXTRA
BASIC FEE		
TOTAL CLAIMS	18 - 20 =	0
INDEP CLAIMS	3 - 3 =	0
____ MULTIPLE DEPENDENT CLAIM PRESENTED		

If the difference in Col. 1 is less than zero, enter "0" in Col. 2.

OTHER THAN A SMALL ENTITY	
RATE	FEE
	\$ 690.00
X \$ 18 =	\$ 0.00
X \$ 78 =	\$ 0.00
+ \$ 260 =	\$ 0.00
TOTAL	\$ 690.00

☒ Please charge my Deposit Account No. 50-0510 in the amount of \$ 690.00.

☒ The Commissioner is hereby authorized to charge payment of the following fees associated with this communication or credit any overpayment to Deposit Account No. 50-0510. A duplicate copy of this sheet is enclosed.

☒ Any additional filing fees required under 37 CFR 1.16.

☒ Any patent application processing fees under 35 CFR 1.17.

Respectfully submitted,

By:

Atty: Charles W. Peterson, Jr.
Registration No. 34,406
Tel. (202) 789-4900
Fax (202) 789-8707
FITCH, EVEN, TABIN & FLANNERY
120 South LaSalle Street
Chicago, Illinois 60603

Certificate of Mailing by Express Mail

"Express Mail" Mailing Label Number
EK754574885US

Date of Deposit 9-29-00

I hereby certify that this paper or fee is being Deposited with the United States Postal Service "Express Mail Post Office to Addressee" Service under 37 CFR Section 1.10 on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, DC 20231

Charles W. Peterson, Jr.

(Typed or printed name of person mailing)

(Signature of person mailing)

NET ZEROING FOR EFFICIENT PARTITION AND DISTRIBUTION

DESCRIPTION

RELATED APPLICATION

5 The present application is related to U.S. Patent Application No. 09/_____
(Attorney Docket No. YOR9-2000-0293-US1) entitled "INDEPENDENT NET TASK
IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION" to
Kimelman et al.; U.S. Patent Application No. 09/_____
10 (Attorney Docket No. YOR9-2000-0464-US1) entitled "MACHINE CUT TASK IDENTIFICATION FOR
EFFICIENT PARTITION AND DISTRIBUTION" to Rajan et al.; and U.S. Patent
Application No. 09/_____
15 (Attorney Docket No. YOR9-2000-0466-US1) entitled
"DOMINANT EDGE IDENTIFICATION FOR EFFICIENT PARTITION AND
DISTRIBUTION" to Wegman et al. all filed coincident herewith and assigned to the
assignee of the present invention.

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention generally relates to distributed processing and more particularly, the present invention relates to efficiently assigning tasks across multiple computers for distributed processing.

Background Description

Any large, multifaceted project, such as a complex computer program, may be segmented into multiple smaller manageable tasks. The tasks then may be distributed amongst a group of individuals for independent completion, e.g., an engineering design project, distributed processing or, the layout of a complex electrical circuit such as a microprocessor. Ideally, the tasks are matched with the skills of the assigned individual and each task is completed with the same effort level as every other task. However, with such an ideal matched task assignment, intertask communication can become a bottleneck to project execution and completion. Thus, to minimize this potential bottleneck, it is important to cluster together individual tasks having the highest level of communication with each other. So, for example, in distributing eight equivalent tasks to pairs of individuals at four locations, (e.g., eight design engineers in four rooms) optimally, pairs of objects or tasks with the highest communication rate with each other are assigned to individual pairs at each of the four locations.

Many state of the art computer applications are, by nature, distributed applications. End-users sit at desktop workstations or employ palmtop information appliances on the run, while the data they need to access resides on distant data servers, perhaps separated from these end-users by a number of network tiers. Transaction processing applications manipulate data spread across multiple servers. Scheduling applications are run on a number of machines that are spread across the companies of a supply chain, etc.

When a large computer program is partitioned or segmented into modular components and the segmented components are distributed over two or more machines, for the above mentioned reasons, component placement can have a significant impact on

program performance. Therefore, efficiently managing distributed programs is a major challenge, especially when components are distributed over a network of remotely connected computers. Further, existing distributed processing management software is based on the assumption that the program installer can best decide how to partition the program and where to assign various program components. However, experience has shown that programmers often do a poor job of partitioning and component assignment.

So, a fundamental problem facing distributed application developers is application partitioning and component or object placement. Since communication cost may be the dominant factor constraining the performance of a distributed program, minimizing inter-system communication is one segmentation and placement objective. Especially when placement involves three or more machines, prior art placement solutions can quickly become unusable, i.e., what is known as NP-hard. Consequently, for technologies such as large application frameworks and code generators that are prevalent in object-oriented programming, programmers currently have little hope of determining effective object placement without some form of automated assistance. En masse inheritance from towering class hierarchies, and generation of expansive object structures leaves programmers with little chance of success in deciding on effective partitioning. This is particularly true since current placement decisions are based solely on the classes that are written to specialize the framework or to augment the generated application.

Furthermore, factors such as fine object granularity, the dynamic nature of object-based systems, object caching, object replication, ubiquitous availability of surrogate system objects on every machine, the use of factory and command patterns, etc., all make partitioning in an object-oriented domain even more difficult. In particular, for conventional graph-based approaches to partitioning distributed applications, fine-grained

object structuring leads to enormous graphs that may render these partitioning approaches impractical.

Finally, although there has been significant progress in developing middleware and in providing mechanisms that permit objects to inter-operate across language and machine boundaries, there continues to be little to help programmers decide object-system placement. Using state of the art management systems, it is relatively straightforward for objects on one machine to invoke methods on objects on another machine as part of a distributed application. However, these state of the art systems provide no help in determining which objects should be placed on which machine in order to achieve acceptable performance. Consequently, the initial performance of distributed object applications often is terribly disappointing. Improving on this initial placement performance is a difficult and time-consuming task.

Accordingly, there is a need for a way of automatically determining the optimal program segmentation and placement of distributed processing components to minimize communication between participating distributed processing machines.

SUMMARY OF THE INVENTION

It is therefore a purpose of the present invention to improve distributed processing performance;

It is another purpose of the present invention to minimize communication between distributed processing machines;

It is yet another purpose of the invention to improve object placement in distributed processing applications;

It is yet another purpose of the invention to determine automatically how objects should best be distributed in distributed processing applications

it is yet another purpose of the invention to minimize communication between objects distributed amongst multiple computers in distributed processing applications.

5 The present invention is a task management system, method and computer program product for determining optimal placement of task components on multiple machines for task execution, particularly for placing program components on multiple computers for distributed processing. First, a communication graph is generated representative of the computer program with each program unit (e.g., an object)
10 represented as a node in the graph. Nodes are connected to other nodes by edges representative of communication between connected nodes. A weight is applied to each edge, the weight being a measure of the level of communication between the connected edges. Terminal nodes representative of ones of the multiple computers are attached to the communication graph. Independent nets may be separated out of the communication
15 graph. For each net, non-terminal nodes adjacent to all of terminal nodes on the net and connected to the net by non-zero weighted edges are identified. For each identified non-terminal node, the smallest weight for any terminal edge is identified and the weight of each terminal edge is reduced by the value of that smallest weight, the weight of terminal edges having the smallest weight being reduced to zero. After reducing weights
20 on any terminal edges, The min cut solution is obtained from the reduced graph and program components, which may be a single program unit or an aggregate of units, are placed on computers according to the communication graph min cut solution.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed preferred embodiment description with reference to the drawings, in which:

5 Figure 1 shows an example of a flow diagram of the preferred embodiment of the present invention wherein a program is segmented, initially, and initial segments are distributed to and executed on multiple computers;

 Figures 2A-C show an example of a communication graph;

10 Figure 3 is a flow diagram of the optimization steps for determining an optimum distribution of program components;

 Figures 4A-B are an example of a simple communication graph reducible using Net Zeroing;

 Figure 5 shows an example of the Net Zeroing steps of the present invention.

15 DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS OF THE INVENTION

As referred to herein, a communication graph is a graphical representation of a multifaceted task such as a computer program. Each facet of the task is an independent task or object that is represented as a node and communication between tasks or nodes is represented by a line (referred to as an edge) between respective communicating nodes.

20 Participating individuals (individuals receiving and executing distributed tasks) are referred to as terminal nodes or machine nodes. A net is a collection of nodes connected together by edges. Two nets are independent if none of the non-terminal nodes of one net shares an edge with a non-terminal node of the other. Thus, for example, a communication graph of a computer program might include a node for each program

object and edges would be between communicating objects, with edges not being included between objects not communicating with each other. In addition, a weight indicative of the level of communication between the nodes may be assigned to each edge. Graphical representation and modeling of computer programs is well known in the art.

Referring now to the drawings, and more particularly, Figure 1 is an example of a flow diagram 100 of the preferred embodiment of the present invention wherein a program is segmented, initially, and initial segments are distributed to and executed on multiple computers. First, in step 102 the communication patterns of a program are analyzed to form a communication graph. Then, in step 104, the traces of the communication graph are analyzed, and an initial partition is determined. In step 106, the partition is optimized for minimum interpartition communication. In step 108, the individual objects are distributed amongst participants for execution according to the optimize partition of step 106.

A component refers to an independent unit of a running program that may be assigned to any participating individual, e.g., computers executing objects of a distributed program. Thus, a component may refer to an instance of an object in an object oriented program, a unit of a program such as Java Bean or Enterprise Java Bean or, a larger collection of program units or components that may be clustered together intentionally and placed on a single participating machine. Further, a program is segmented or partitioned into segments that each may be a single component or a collection of components. After segmenting, analyzing the segments and assigning each of segments or components to one of the multiple participating machines or computers according to the present invention, the final machine assignment is the optimal assignment.

Thus, a typical communication graph includes multiple nodes representative of components with weighted edges between communicating nodes. Determining communication between components during program execution may be done using a typical well known tool available for such determination. Appropriate communication determination tools include, for example, Jinsight, for Java applications that run on a single JVM, the Object Level Tracing (OLT) tool, for WebSphere applications or, the monitoring tool in Visual Age Generator.

Figures 2A-C show an example of a communication graph of a net 110 that includes multiple nodes 112, 114, 116, 118, 120 and 122. Each node 112, 114, 116, 118, 120 and 122 represents a program component connected to communication edges 124, 126, 128, 130, 132, 134, 136 and 138 to form the net 110. Adjacent nodes are nodes that share a common edge, e.g., nodes 114 and 122 share edge 126. Each edge 124, 126, 128, 130, 132, 134, 136 and 138 has been assigned a weight proportional to, for example, the number of messages between the adjacent components.

In Figure 2B, Machine nodes 140, 142 and 144 representative of each participating machine (three in this example) are shown connected by edges 146, 148, 150. Initially, a node 112, 114, 116, 118, 120 and 122 may be placed on a machine 140, 142, 144 by adding an edge 146, 148, 150 with infinite weight (indicating constant communication) between the node and the machine. Typically, initial assignment places nodes with specific functions (e.g., database management) on a machine suited for that function. After the initial placement assigning some nodes 112, 114 and 122 to machines 140, 142, 144, other nodes 116, 118, 120 are assigned to machines 140, 142, 144, if they communicate heavily with a node 112, 114, 122 already assigned to that machine 140, 142, 144. Additional assignment is effected by selectively collapsing edges, combining the nodes on either end of the collapsed edge and re-assigning edges that were attached to

one of the two former adjacent nodes to the combined node. When assignment is complete, all of the nodes 112, 114, 116, 118, 120 and 122 will have been placed on one of the machines at terminal nodes 140, 142, 144 and the final communication graph may be represented as terminal nodes 140, 142, 144 connected together by communication edges.

For this subsequent assignment, the graph is segmented by cutting edges and assigning nodes to machines as represented by 152, 154 and 156 in Figure 2C to achieve what is known in the art as a minimal cut set or min cut set. A cut set is a set of edges that, if removed, eliminate every path between a pair of terminal nodes (machine nodes) in the graph. A min cut set is a cut set wherein the sum of the weights of the cut set edges is minimum. While there may be more than one min cut set, the sum is identical for all min cut sets. A min cut may be represented as a line intersecting the edges of a min cut set. So, in the example of Fig. 2C, the sum of the weights of edges 124, 126, 128, 132 and 138 is 2090, which is cost of the cut and is representative of the total number of messages that would be sent between machines at terminal nodes 140, 142, 144 with this particular placement. The min cut identifies the optimum component placement with respect to component communication. While selecting a min cut set may be relatively easy for this simple example, it is known to increase in difficulty exponentially with the number of nodes in the graph.

Figure 3 is a flow diagram 160 of the optimization steps for determining an optimum distribution of program components to individual participating computers according to a preferred embodiment of the present invention. First, in step 162, an initial communication graph is generated for the program. Then, in step 164 machine nodes are added to the communication graph. As noted above, certain types of components are designated, naturally, for specific host machine types, e.g., graphics components are

designated for clients with graphics capability or, server components designated for a data base server. After assigning these host specific components, in step 168 independent nets are identified and the communication graph is partitioned into the identified independent nets as described in U.S. Patent Application No. 09/_____ (Attorney Docket No. YOR9-2000-0293-US1) entitled "INDEPENDENT NET TASK IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION" to Kimelman et al., assigned to the assignee of the present invention and incorporated herein by reference. In step 170 a min cut solution is derived for each of the independent nets using the preferred embodiment Net Zeroing reduction method described hereinbelow to reduce the independent nets. The min cuts for all of the independent nets being the min cut for the whole communication graph.

Figures 4A-B are an example of a simple communication graph 180 reducible using the preferred embodiment Net Zeroing method of the present invention. In this example, the graph 180 includes four (4) non-terminal nodes 182, 184, 186 and 188 connected together by edges 190, 192 and 194, referred to herein as non-terminal edges. Three (3) terminal nodes 196, 198 and 200 are connected to non-terminal nodes 184, 186 and 188 by edges 202, 204, 206, 208, 210 and 212. Non-terminal node 182 is connected to all three terminal nodes 196, 198 and 200 by edges 214, 216 and 218, respectively. Edges 202-214 are referred to herein as terminal edges.

A weight is represented as being attached to each edge 190 - 194 and 202 - 218. Essentially, Net Zeroing is applied to any node in a net that is adjacent (has a weight greater than zero) to all terminal nodes in the net, i.e., node 182 in this example. Net Zeroing reduces the weight of each terminal node attached to such a node by a value equal to the smallest terminal node weight. In this example, terminal edge 216 and 218 have the smallest terminal edge weight (4) of the terminal edges 214, 216 and 218

attached to node 182. So, in Figure 4B, the weight attached to terminal edges 214, 216 and 218 is shown with each reduced by four. This decrease in edge weights affects all possible cuts equally, and thus has no effect on which edges are present in the multiway minimum cut solution. It does however, result in a net that is more amenable to finding the min cut solution (using, for example, the Dominant Edge heuristic). Thus, the graph has been reduced in the sense that it is has been simplified, making the min cut solution easier to find, much more quickly.

Figure 5 is an example of the Net Zeroing steps 220 of identifying terminal edges wherein weights are reduced according to the preferred embodiment of the present invention. First, in step 222, an independent net is checked to determine whether it includes any non-terminal nodes that are adjacent to all machine nodes connected to that net. Net Zeroing may be applied to any such identified non-terminal node wherein the weight of all connected non-terminal edges is greater than zero. It should be noted that the weight of one or more terminal edges may be zero as a result of a prior Net Zeroing application as in the example of Figure 4B. If no non-terminal nodes are found adjacent to all terminals and connected thereto by terminal edges with weights that are all greater than zero, then in step 224 Net Zeroing is determined to be inapplicable to the net. Otherwise, in step 226, the minimum weight for the terminal edges is identified. In step 228, the weight of each of the non-terminal edges is reduced by the value of the minimum edge weight. It should be noted that non-terminal edge weights remain unchanged after step 228. In step 230, an indication is made that Net Zeroing was successful. Finally, in step 242, the reduced net is returned for determination of the min cut solution.

In other words, for each nonterminal node 182 connected to all terminal nodes 196, 198 and 200, let c be the smallest amount of communication represented by edges 216, 218 between an un-hosted component A on non-terminal node 182 and components

hosted on a machine at terminal nodes 198 or 200, where the minimum is taken over the set of machines that communicates with the components in the independent net 180 containing A at node 182. Then, subtracting c from the amount of communication between A at node 182 and each of the machines at terminal nodes 196, 198, 200 in the set does not change the min cut solution, but simplifies arriving at that solution.

So, the preferred embodiment, the min cut step 170 is used iteratively, wherein Net Zeroing as described herein is used to reduce independent nets in combination with the Machine Cut reduction method of U.S. Patent Application No. 09/_____ (Attorney Docket No. YOR9-2000-0464-US1) entitled "MACHINE CUT TASK IDENTIFICATION FOR EFFICIENT PARTITION AND DISTRIBUTION" to Rajan et al., and the Dominant Edge method of U.S. Patent Application No. 09/_____ (Attorney Docket No. YOR9-2000-0466-US1) entitled "DOMINANT EDGE FOR EFFICIENT PARTITION AND DISTRIBUTION" to Roth et al., both filed coincident herewith, assigned to the assignee of the present invention and incorporated herein by reference. Further, as independent nets are reduced, those reduced nets are further checked as in step 168 above to determine if they may be divided into simpler independent nets. Then, the Net Zeroing method of the preferred embodiment is again applied to those simpler independent nets, along with the Machine Cut method and Dominant Edge method, if necessary. To reach a solution more quickly, on each subsequent pass, only nodes and edges of a subgraph that were adjacent to areas reduced previously are rechecked. Thus, the communication graph is simplified by eliminating edges to reach a min cut solution much quicker and much more efficiently than with prior art methods.

The reduction method of the preferred embodiment reduces the number of independent components in the communication graph of a complex program. In the best

case, an appropriate allocation of every component in the program is provided. However, even when best case is not achieved, the preferred embodiment method may be combined with other algorithms and heuristics such as the branch and bound algorithm or the Kernighan-Lin heuristic to significantly enhance program performance. Experimentally, the present invention has been applied to communication graphs of components in several programs with results that show significant program allocation improvement, both in the quality of the final solution obtained and in the speed in reaching the result.

Although the preferred embodiments are described hereinabove with respect to distributed processing, it is intended that the present invention may be applied to any multi-task project without departing from the spirit or scope of the invention. Thus, for example, the task partitioning and distribution method of the present invention may be applied to VLSI design layout and floor planning, network reliability determination, web pages information relationship identification, and "divide and conquer" combinatorial problem solution approaches, e.g., "the Traveling Salesman Problem."

While the invention has been described in terms of preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.

CLAIMS

What is claimed is:

1. A task management method for determining optimal placement of task components, said method comprising:

5 a) generating a communication graph representative of a task, task components represented as nodes of said communication graph and edges connecting ones of said nodes, said edges representing communication between connected nodes and being weighted proportional to communication between connected nodes;

b) assigning terminal nodes to said communication graph;

10 c) identifying nodes adjacent to terminal nodes and connected to each adjacent terminal node by a terminal edge;

d) reducing the weight of each terminal edge for each said identified node by the minimum weight of every terminal edge for said identified node;

e) determining a min cut solution for said communication graph; and

15 f) placing task components on said terminal nodes responsive to said min cut solution.

2. A task management method as in claim 1, after the step (b) of assigning terminal nodes, further comprising the step of:

20 b1) identifying independent nets in said communication graph, each of said independent nets being connected between a plurality of said terminal nodes.

3. A task management method as in claim 2, wherein the step (c) of identifying nodes comprises the steps of:

i) identifying nodes that are adjacent to each of said plurality of terminal nodes;

5 ii) selecting nodes from said identified nodes, terminal edges connected to said selected nodes having a weight greater than zero;

iii) identifying the minimum terminal edge weight for each said selected node; and

10 iv) reducing weights of all terminal edges of each selected node by its corresponding identified minimum.

4. A task management method as in claim 3, wherein each said task component is a unit of the computer program.

5. A task management method as in claim 4, wherein said each computer program unit is an instance of an object in an object oriented program

15 6. A task management method as in claim 4, wherein in step (d) computer program units are placed on computers, computer program units being placed on a common computer being combined into a single component.

20 7. A task management method as in claim 3 wherein said task is integrated circuit chip functional element placement and said task components are logic elements, said logic elements being placed on an integrated circuit chip in placement step (e).

8. A distributed processing system for determining optimal placement of computer program components on multiple computers, said distributed processing system comprising:

means for generating a communication graph of nodes interconnected by edges and representative of a computer program, computers executing said computer program being represented as terminal nodes, computer program components being represented as non-terminal nodes and edges connecting ones of said non-terminal nodes, said edges representing communication between connected nodes and being weighted proportional to communication between connected nodes;

means for identifying non-terminal nodes connected to terminal nodes;
means for identifying a minimum terminal edge weight;
means for reducing terminal edge weights responsive to said identified minimum weight;

means for determining a min cut solution for said communication graph; and
means for placing program components on ones of said computers responsive to said determined min cut solution;
said computer program being executed by said computers.

9. A distributed processing system as in claim 8, further comprising:
means for identifying independent nets connected between a plurality of said terminal nodes.

10. A distributed processing system as in claim 9 wherein said means for identifying non-terminal nodes connected to terminal nodes identifies non-terminal nodes connected to all terminal nodes connected to an independent net.

11. A distributed processing system as in claim 10, wherein each said program component is a unit of the computer program.

12. A distributed processing system as in claim 11, wherein said each program unit is an instance of an object in an object oriented program

5 13. A computer program product for determining optimal placement of functional components, said computer program product comprising a computer usable medium having computer readable program code thereon, said computer readable program code comprising:

10 computer readable program code means for generating a communication graph of nodes interconnected by edges and representative of a function, said nodes including a plurality of terminal nodes, functional components being represented as non-terminal nodes, said edges representing communication between connected nodes;

computer readable program code means for identifying non-terminal nodes connected to terminal nodes;

15 computer readable program code means for identifying a minimum terminal edge weight;

computer readable program code means for reducing terminal edge weights responsive to said identified minimum weight;

20 computer readable program code means for determining a min cut solution for said communication graph; and

computer readable program code means for placing program components on said terminal nodes responsive to said determined min cut solution.

14. A computer program product as in claim 13, further comprising:
computer readable program code means for identifying independent nets
connected between a plurality of said terminal nodes.

5 15. A computer program product as in claim 14, wherein the computer readable
program code means for identifying non-terminal nodes connected to terminal nodes
identifies non-terminal nodes connected to all terminal nodes connected to an
independent net.

16. A computer program product as in claim 15, wherein said function is a computer
program and each said functional component is a unit of the computer program.

10 17. A computer program product as in claim 16, wherein each said program unit is an
instance of an object in an object oriented program.

18. A computer program product as in claim 15, wherein said function is an integrated
circuit and each said functional component is a logic element.

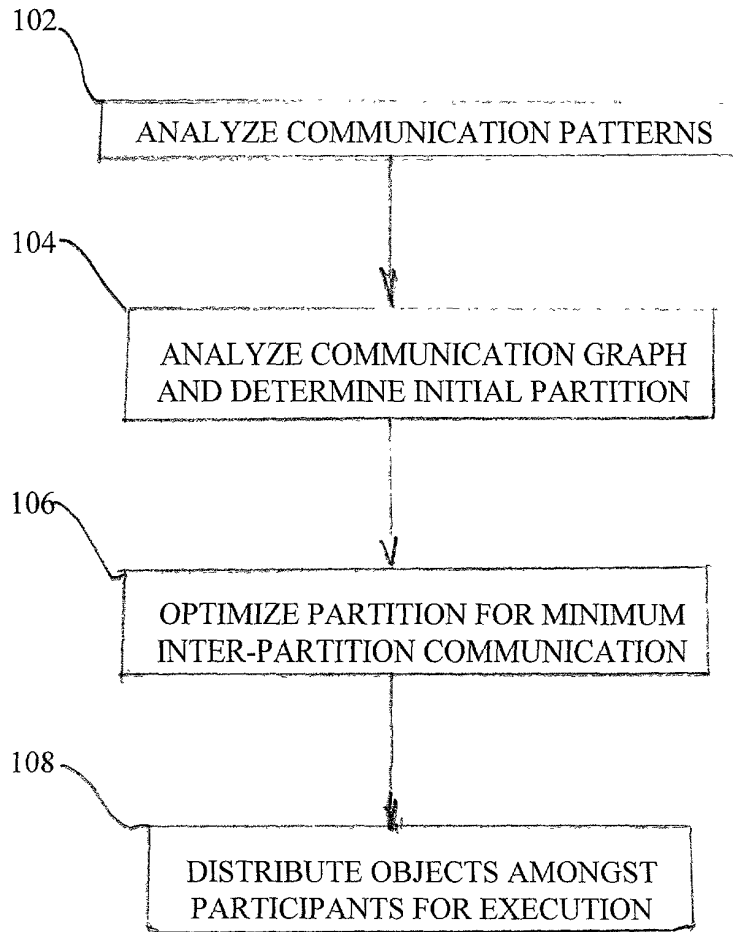


FIG. 1

006260" 52792960




Fig. A

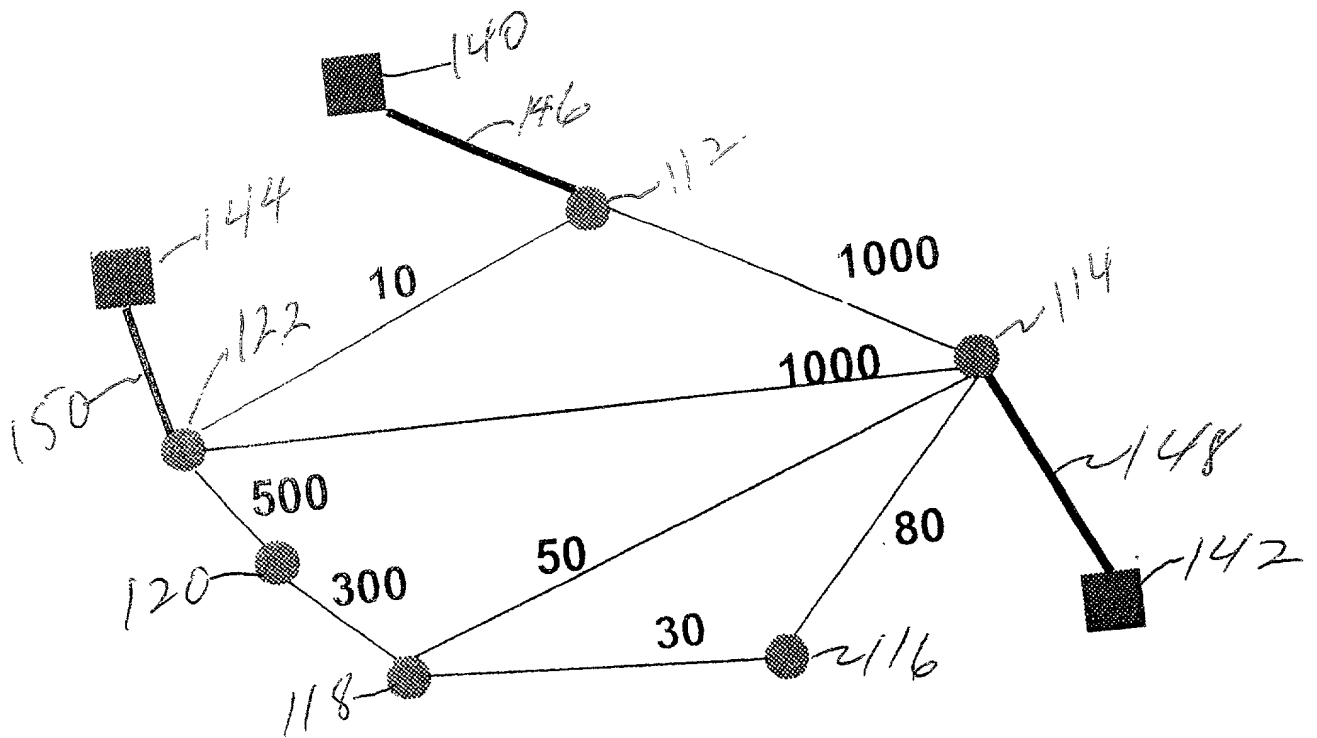


Fig. 2B

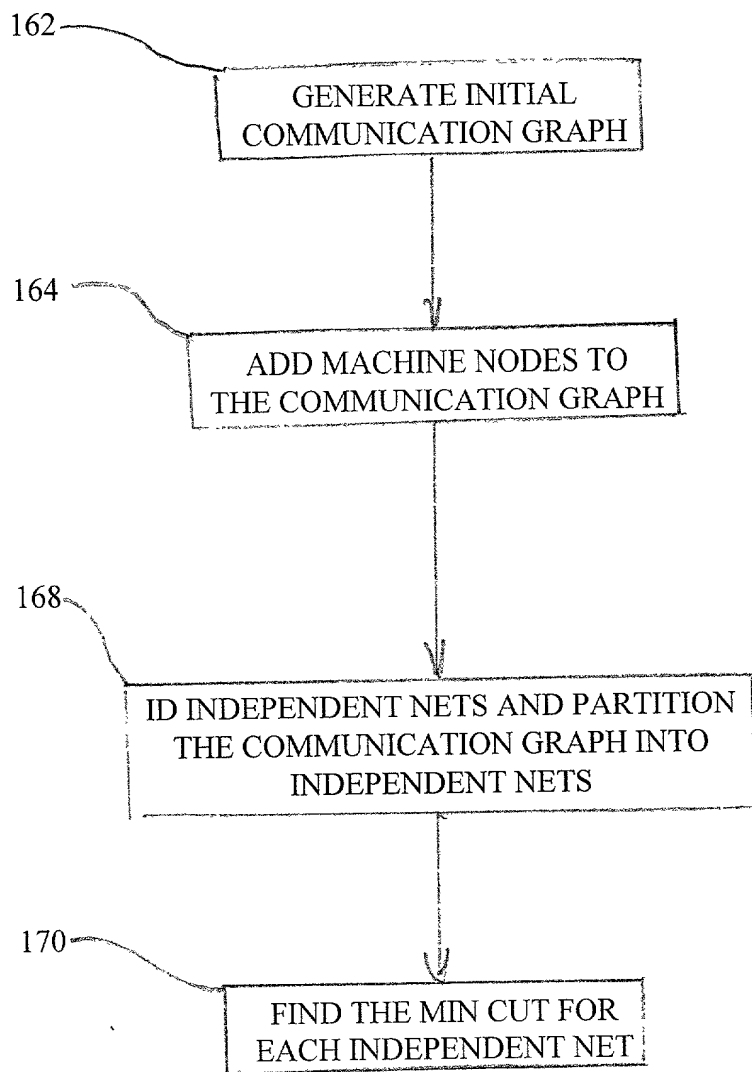


FIG. 3

Fig. 4A

006360" 00732360


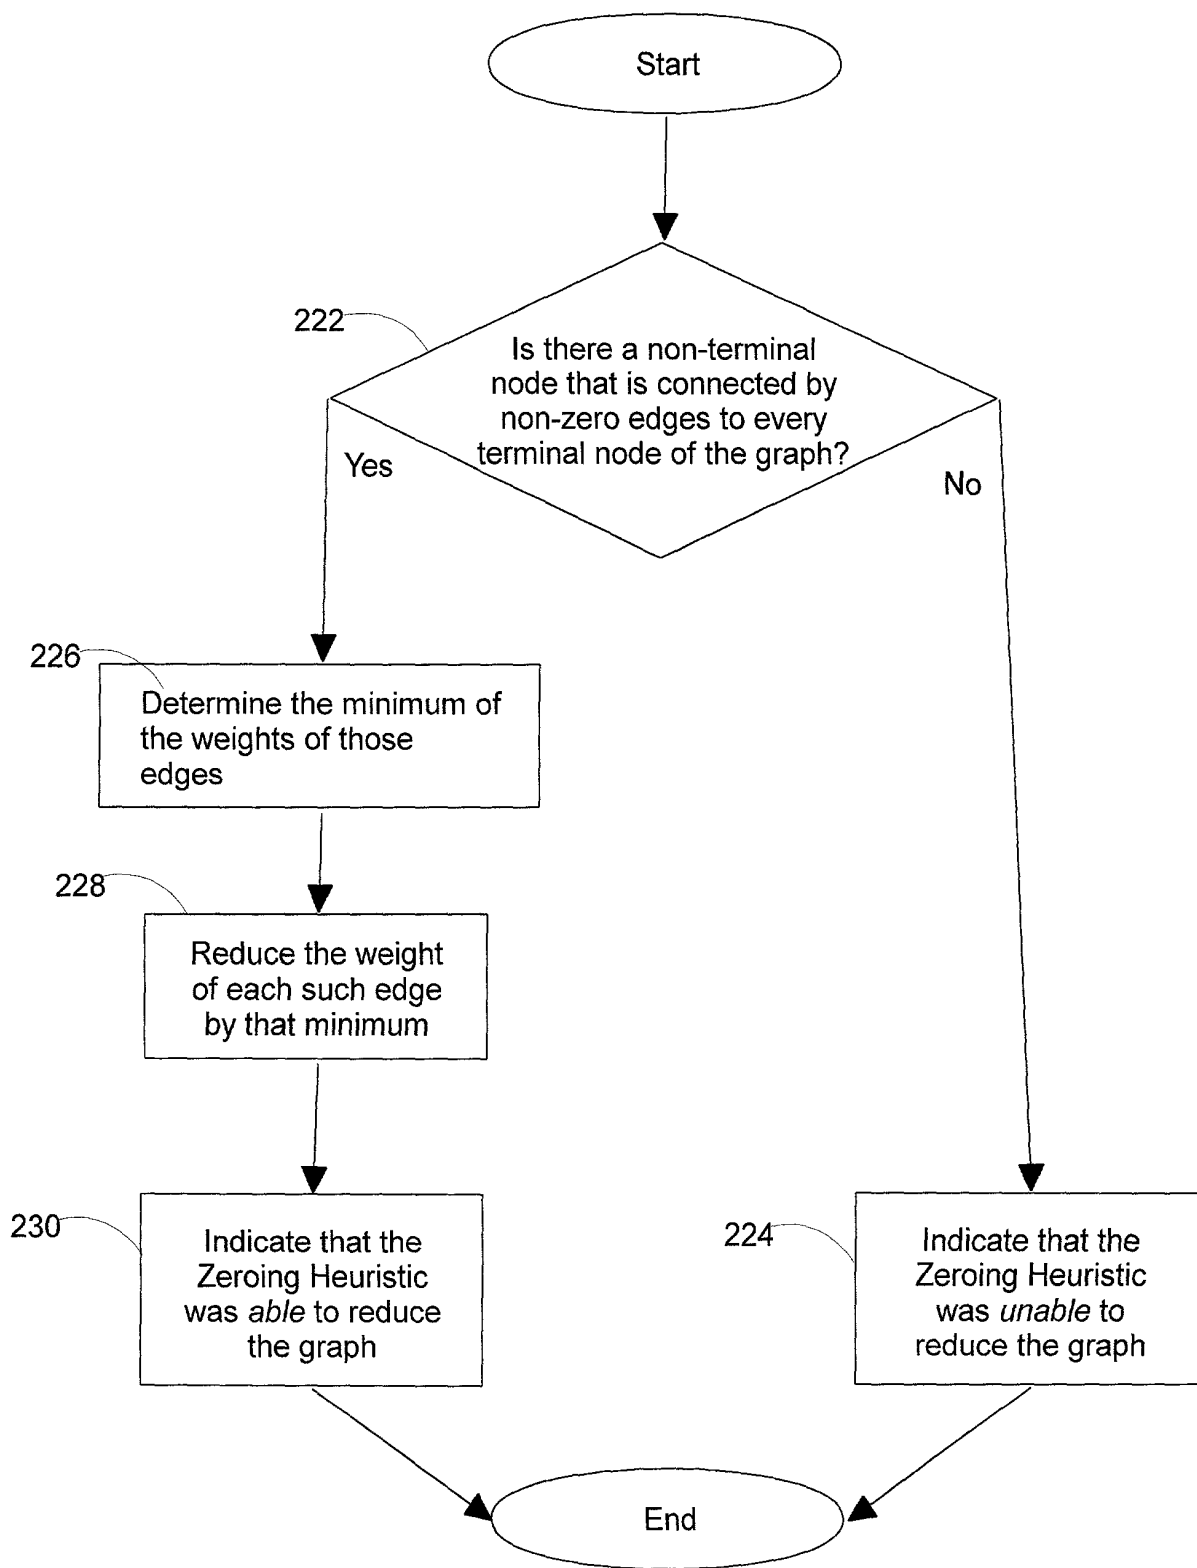
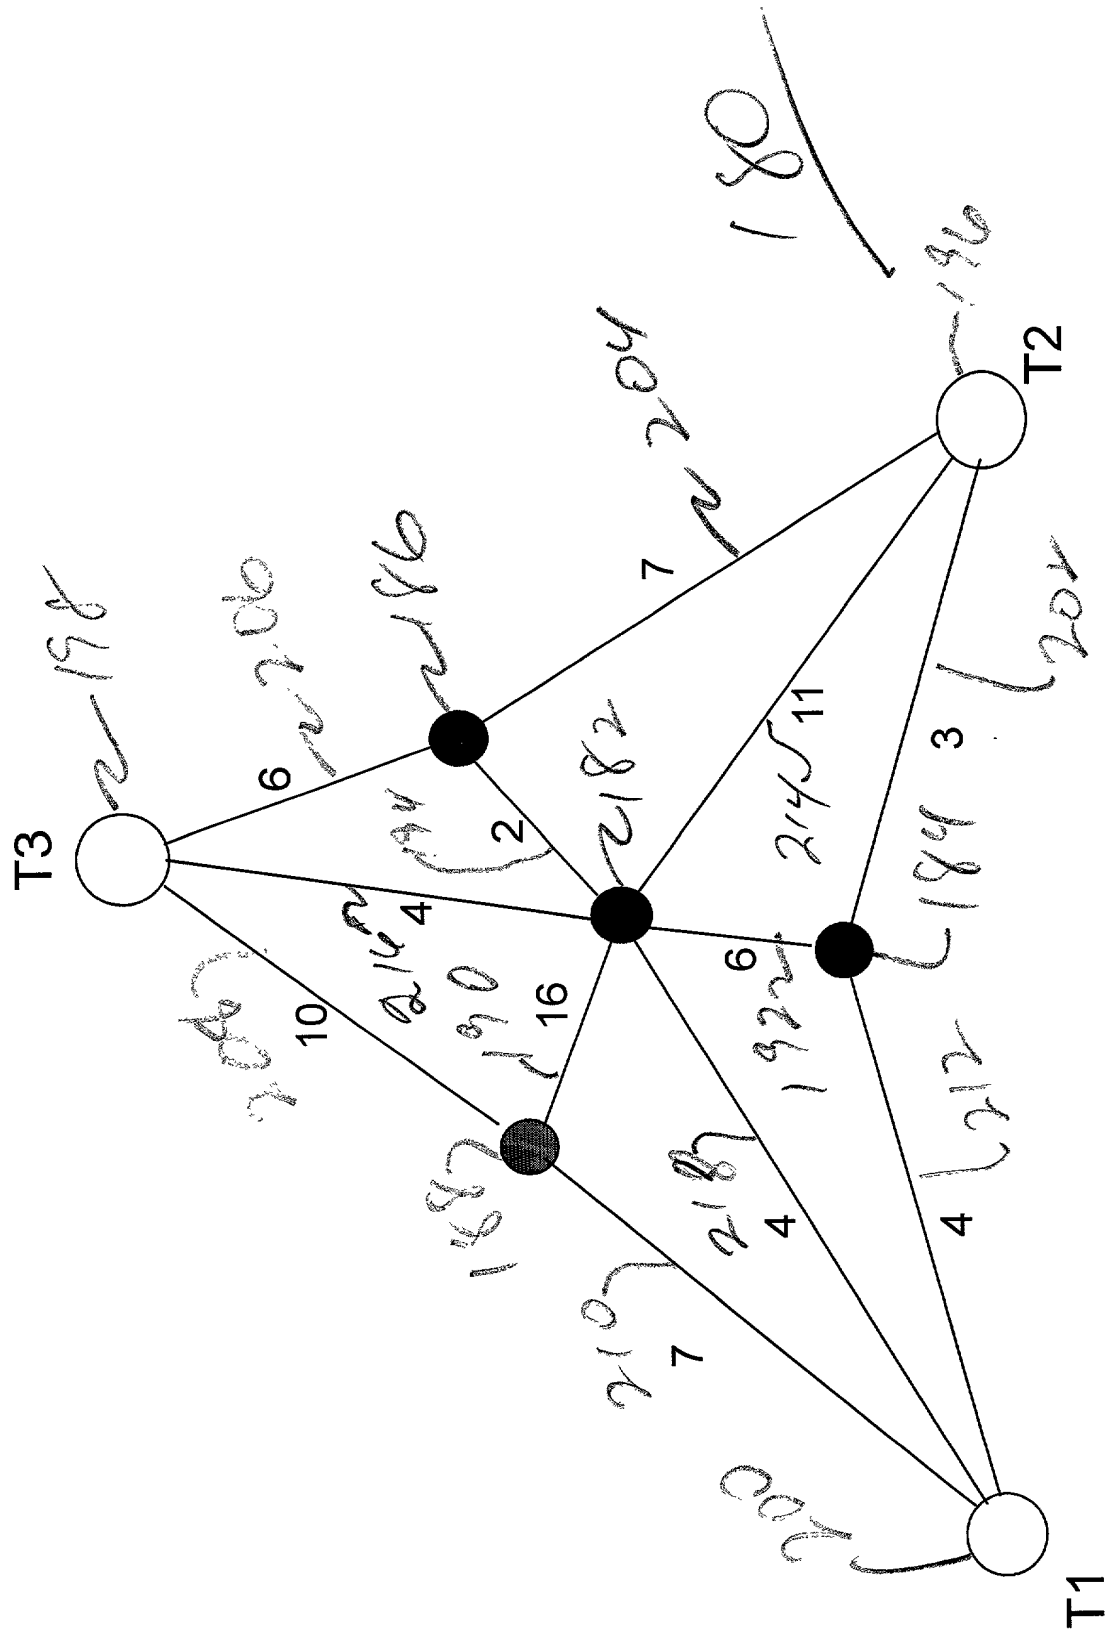


Fig. 4B



220

Fig. 5



Initial Graph for zeroing example

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

NET ZEROING FOR EFFICIENT PARTITION AND DISTRIBUTION

the specification of which (check one)

X is attached hereto.

_____ was filed on _____ as United States Application Number

or PCT International Application Number _____

and was amended on _____ (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119(a)-(d) or §365(b) of any foreign application(s) for patent or inventor's certificate, or §365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate, or PCT International application, having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)			Priority Claimed	
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	____ Yes ____ No	
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	____ Yes ____ No	
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	____ Yes ____ No	

I hereby claim the benefit under 35 U.S.C. §119(e) of any United States provisional application(s) listed below.

_____ (Application Number)	_____ (Filing Date)
_____ (Application Number)	_____ (Filing Date)

I hereby claim the benefit under 35 U.S.C. §120 of any United States Application(s), or §365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States, or PCT International application in the manner provided by the first paragraph of 35 U.S.C. §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in 37 CFR §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

_____ (Application Serial No.)	_____ (Filing Date)	_____ (Status) (patented, pending, abandoned)
_____ (Application Serial No.)	_____ (Filing Date)	_____ (Status) (patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith (list name and registration number).

Manny W. Schecter (Reg. 31,722), Lauren C. Bruzzone (Reg. 35,082), Christopher A. Hughes (Reg. 26,914), John E. Hoel (Reg. 26,279), Joseph C. Redmond, Jr. (Reg. 18,753), Stephen C. Kaufman (Reg. 29,551), Jay P. Sbrollini (Reg. 36,266), David M. Shofi (Reg. 39,835), Robert M. Trepp (Reg. 25,993), Louis P. Herzberg (Reg. 41,500), Daniel P. Morris (Reg. 32,053), Paul J. Otterstedt (Reg. 37,411), Louis J. Percello (Reg. 33,206), Douglas W. Cameron (Reg. 31,596), Wayne L. Ellenbogen (Reg. 43,602) and Marian Underweiser (Reg. 46,134)

Send Correspondence to: Charles W. Peterson, Jr. Fitch, Even, Tabin & Flannery

120 South LaSalle Street, Suite 1600, Chicago, IL 60603

Direct Telephone Calls to: (name and telephone number) Charles W. Peterson, (202) 789-4900

Tova Roth

Full name of sole or first inventor

Inventor's Signature

Date

408 Edward Avenue, Woodmere, NY 11598

Residence

USA

Citizenship

Same as above

Post Office Address

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATIONVadakkedathu T. RAJAN**Full name of second joint-inventor, if any**

Inventor's signature

Date

223 Schrade Road, Briarcliff Manor, NY 10510
ResidenceUSA

Citizenship

same as above

Post Office Address

Douglas N. KIMELMAN**Full name of third joint-inventor, if any**

Inventor's signature

Date

P.O. Box 704, Yorktown Heights, NY 10598
ResidenceCanada

Citizenship

Same as above

Post Office Address

Mark N. WEGMAN**Full name of fourth joint-inventor, if any**

Inventor's Signature

Date

43 Overlook Road, Ossining, NY 10562
ResidenceUSA

Citizenship

Same as above

Post Office Address

Karin HOGSTEDT**Full name of fifth joint inventor, if any**

Inventor's Signature

Date

Room D275, 180 Park Avenue, Florham Park, NJ 07932
ResidenceSweden

Citizenship

Same as above

Post Office Address

Full name of sixth joint-inventor, if any

Inventor's signature

Date

Residence

Citizenship

Post Office Address